

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 11 Aug 2016		2. REPORT TYPE Journal Article		3. DATES COVERED (From – To) Feb 2016 – May 2016	
4. TITLE AND SUBTITLE PERCEPTUAL PERFORMANCE IMPACT OF GPU-BASED WARP & ANTI-ALIASING FOR IMAGE GENERATORS				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Logan Williams, Charles Bullock, Doug Traill, Kevin Amaratunga, James Gaska, Marc Winterbottom				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) USAF School of Aerospace Medicine Aeromedical Research Department/FHOH 2510 Fifth St., Bldg. 840 Wright-Patterson AFB, OH 45433-7913				8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-SA-WP-JA-2016-0028	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSORING/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.					
13. SUPPLEMENTARY NOTES Cleared, 88PA, Case # 2016-2978, 10 Jun 2016. IMAGE 2016 Conference Dayton, Ohio – 28-29 June 2016					
14. ABSTRACT In 2012 the U.S. Air Force School of Aerospace Medicine, in partnership with the Air Force Research Laboratory (AFRL) and NASA AMES, constructed the Operational Based Vision Assessment (OBVA) simulator. This 15-channel, 150-megapixel display system remains one of the highest resolution displays ever built. One of the original goals for the simulator was to implement a distortion correction system that introduces “zero” frame latency into the overall system. This distortion correction was achieved using a combination of Scalable Display’s EasyBlend SDK and NVIDIA’s Warp and Intensity adjustment API. This paper describes the results of a collaboration between USAFSAM, Scalable, and NVIDIA to evaluate NVIDIA’s WARP 2.0 API, which allows for several user-selectable filtering techniques. These filters have the potential to improve the quality of the display warp and improve anti-aliasing performance without change to the low latencies already achieved. This paper provides a brief review of the different filtering techniques under investigation, as well as an assessment of their performance within a flight simulation environment. The evaluation has been conducted using psychometric methods to determine threshold performance of human observers on an operationally relevant aircraft orientation task conducted at an eye-limiting resolution (1 arcmin/lp).					
15. SUBJECT TERMS Antialiasing, image generator, perception, geometric distortion, geometric warp, NVIDIA warp API, GPU warp, vision assessment.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON Logan Williams
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code)

PERCEPTUAL PERFORMANCE IMPACT OF GPU-BASED WARP & ANTI-ALIASING FOR IMAGE GENERATORS

Logan Williams¹, Charles Bullock¹, Doug Traill², Kevin Amaratunga³, James Gaska¹, Marc Winterbottom¹

¹OBVA Laboratory, U.S. Air Force School of Aerospace Medicine, Wright-Patterson AFB, OH

²NVIDIA Corp.

³Scalable Display Technology, Inc.

ABSTRACT

In 2012 the U.S. Air Force School of Aerospace Medicine, in partnership with the Air Force Research Laboratory (AFRL) and NASA AMES, constructed the Operational Based Vision Assessment (OBVA) simulator. This 15-channel, 150-megapixel display system remains one of the highest resolution displays ever built. One of the original goals for the simulator was to implement a distortion correction system that introduces “zero” frame latency into the overall system. This distortion correction was achieved using a combination of Scalable Display’s EasyBlend SDK and NVIDIA’s Warp and Intensity adjustment API. This paper describes the results of a collaboration between USAFSAM, Scalable, and NVIDIA to evaluate NVIDIA’s WARP 2.0 API, which allows for several user-selectable filtering techniques. These filters have the potential to improve the quality of the display warp and improve anti-aliasing performance without change to the low latencies already achieved. This paper provides a brief review of the different filtering techniques under investigation, as well as an assessment of their performance within a flight simulation environment. The evaluation has been conducted using psychometric methods to determine threshold performance of human observers on an operationally relevant aircraft orientation task conducted at an eye-limiting resolution (1 arcmin/lp).

INTRODUCTION

In 2012 the U.S. Air Force School of Aerospace Medicine, in partnership with the Air Force Research Laboratory (AFRL) and NASA AMES, constructed the Operational Based Vision Assessment (OBVA) simulator to evaluate the relationship between clinical/laboratory measures of ocular health and operational task performance (Figure 1) [1]. To support this line of research, a front projection dome geometry flight simulator was developed with a wide 160° x 60° field of view. To approximate optical infinity, a 4-meter eye-relief was chosen, which necessitates the use of

multiple high-resolution projectors to achieve eye-limiting resolution [2]. In total, 15 Barco/Esterline Sim 10 LCOS projectors were employed to achieve a resolution of 1 arcmin per line pair (0.5 arcmin/pixel).

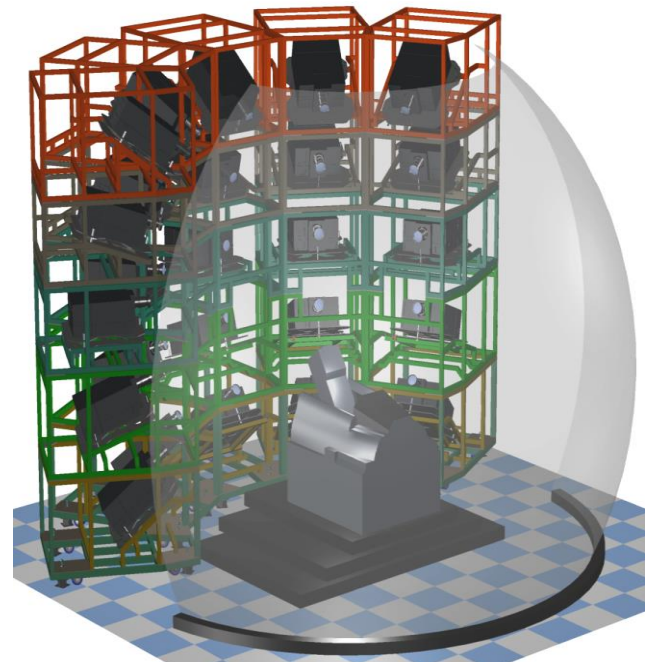


Figure 1: Schematic illustration of the OBVA high-bay simulation visual display system.

BACKGROUND

The implementation of the OBVA simulator necessitated a multi-projector design in which off-axis projectors are used to display images on a spherical surface. Therefore, a method of geometry correction and luminance blending between projector channels was required to unite the individual projections into a single, apparently seamless image. The Scalable Display Manager software, developed by Scalable Display Technologies, Inc., was selected to provide geometry correction using bilinear interpolation.

This software was implemented using NVIDIA Quadroplex 6000 graphics cards. However, since the initial build of the OBVA simulator, NVIDIA has implemented several more advanced geometric warp features in their DesignWorks SDK for Quadro GPUs. Working with Scalable Display Technologies to implement these new geometric warp features, the OBVA lab conducted a perceptual performance evaluation, using psychometric methods, to quantify the level of improvement these alternate warp methods might impart upon high-acuity human observers using the OBVA simulator. The goal was to determine if an optimal geometric warp method may improve eye-limited human performance.

FILTERING METHODS

The NVIDIA Warp API (i.e., Warp 2.0) provides support for categories of operations that range beyond the scope of those found in familiar graphics APIs such as DirectX and OpenGL and were specifically developed to support applications that require projector blending, projection mapping, keystone correction, as well general purpose desktop remapping tool. The original API used bilinear filtering, which provided a very fast method for warping while maintaining sharpness of the image for detailed items such as text. However, for some applications this may introduce aliasing artifacts. The new Warp 2.0 methods allow the developer to choose the correct filtering method based on the display content. The new Warp API contains seven developer-selectable warp methods, consisting of four fixed filters and three adaptive filters:

Fixed filters:

- 1) BiLinear resampling (original method)
- 2) BiCubic Triangular resampling
- 3) BiCubic Bell function resampling
- 4) BiCubic Bspline resampling

Adaptive filters:

- 5) Adaptive BiCubic Triangular resampling
- 6) Adaptive Bell function resampling
- 7) Adaptive Bspline resampling

The adaptive filters differ from the fixed filters in that they apply a certain degree of low-pass filtering so as not to blur the image in regions that require little or no warping. Note that in each case the original warp mesh remains unaltered, only the interpolation method of the filter is changed.

Computational Performance

The APIs are applied as part of the scan-out process, independent of an application. This creates a low-latency path for warping the desktop that doesn't involve any off-screen rendering or texture manipulation. This is especially useful for very high-resolution displays where any off-screen copies will likely result in additional latency.

The original bilinear filtering method introduced a 0.1-ms overhead to the render process, which was verified in the OBVA simulator using QuadroPlex 6000 cards based upon the older Fermi GPU architecture. Because the newer warp methods are designed to operate on newer GPU architectures (i.e., Kepler or Maxwell), rather than the older Fermi architecture, a significant decrease in performance was observed, with frame rates dropping to approximately one-half to one-third of the intended 60 Hz. However, using modern GPU architectures, the new warp methods introduce minimal latency (~0.4 to ~1.0 ms) as shown in Figure 2.

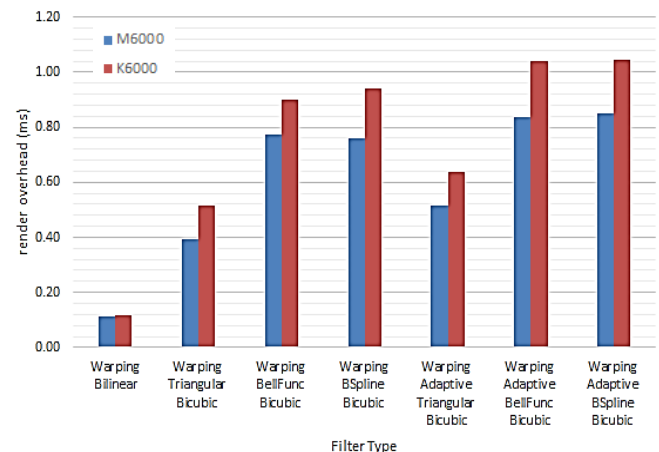


Figure 2: Comparison of increased GPU latency for each of seven warp methods using a 1 to 1 pixel mapping. Note, these comparisons are approximate and will vary slightly with different applications.

VISUAL ASSESSMENT

The possibility of minimizing aliasing artifacts using these additional warp methods is of primary interest in flight simulation, in which target identification at great distances is a critical performance task. Small targets, in any visual display system, will exhibit significant aliasing when the target size approaches a few pixels. Although such aliasing is unavoidable, it is desirable that no additional aliasing is introduced for images undergoing geometric correction (i.e., warp).

To assess the impact of aliasing, due to geometric warp, on

operational performance, the OBVA lab conducted a psychometric evaluation of target aspect/orientation using each of the seven available warp methods. Human observers with very high acuity, near the display limit of the OBVA simulator, were specifically chosen to better evaluate the effects of subtle (e.g. single pixel) changes due to interpolation method which may be imperceptible to normal-acuity or lower-acuity observers. This task was repeated in three regions of the visual display, which were specifically chosen due to the different degrees of local geometric warp. The center of the projector was chosen for minimum localized warp, while a second region near the left edge (but not in a blend zone) was used to represent a region with greater localized warp. A third region was evaluated in the middle of a horizontal blend zone to assess any possible effect due to image blending. These regions are shown in Figure 3.



Figure 3: Projector regions for psychometric evaluation. Position 1 is centered on one projector (0° offset). Position 2 is near the edge (10° offset). Position 3 is located in the center of a blend zone between two projectors (15.2° offset).

Operational Task

The operational task is similar to those used by USAFSAM for past image generator research and consisted of an aspect/roll identification task, in which a single aircraft model was rolled $\pm 20^\circ$ from horizontal on each trial, with 50 trials per condition [3]. The rolling action of the model was animated such that the roll action spanned 5 frames. The aircraft remained in the $\pm 20^\circ$ position for 2 seconds prior to rolling back to horizontal. The user was required to push a button indicating the direction of the roll and feedback, in the form of a correct response or incorrect response tone, was provided after each response. Although there was no time limit placed on subject responses, the stimulus ($\pm 20^\circ$ roll) was only present for 2 seconds.

The well-accepted Ψ (psi) method [4, 5] was adopted, which uses prior observer responses to optimize the information gain on the next trial to estimate a psychometric function. A typical Ψ estimate of the psychometric function is shown in Figure 4 (top). The Ψ algorithm sampled

distance using 0.05 log unit steps, such that the distance between the ownship and target model is $10^{\Delta \text{stim level}}$ in meters. Typical stim levels ranged from 3.5 to 4.0, which correspond to 3.2 to 10 km. The red curve is the maximum likelihood fit of a Weibull psychometric function that relates distance to the proportion of correct responses. Distance threshold was taken as the distance corresponding to a 0.81 proportion correct.

In Figure 4 (top) the size of the data points are proportional to the square root of the number of trials at that distance and it can be seen that majority of trials used stimuli that fall in the information rich zone between perfect performance and chance performance. The bottom chart in Figure 4 shows stimulus distance on each trial and again shows how the psi method converges to distances near threshold levels.

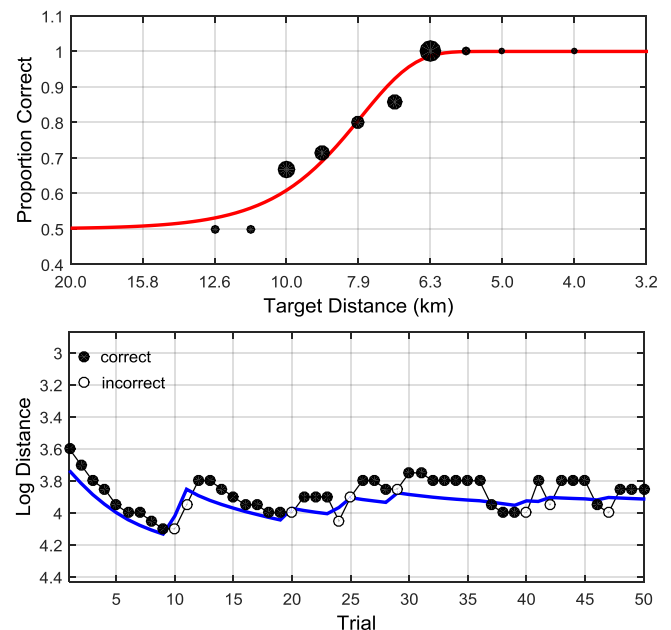


Figure 4. Typical estimate of the psychometric function given by the Ψ method (top) and convergence action of the adaptive algorithm (bottom).

Objective Measures

In addition to psychometric measures of human performance, several camera-based luminance measurements were conducted to determine the contrast ratio of various target features near threshold for each region of the display using a Radiant Prometric imaging colorimeter (Model PM-1433F-1). In particular, the contrast ($[\text{target-background}]/\text{background}$) of both the fuselage and wingtips relative to the background were measured. Figure 5 illustrates the visual difference

between bilinear and bicubic triangular warp methods for an aircraft at 1 km (stim level = 3.0) in the center position. These two warp methods produced the greatest apparent visual difference, with the alternate methods falling somewhere in-between. Figure 6 details the level of aliasing present for bilinear warp in the central position at various stim levels. At small target sizes, the wings are the first model feature to experience severe aliasing and are essentially removed from the target at distances greater than 6 km (stim level ~3.8). Beyond this distance, the fuselage provides the dominant aspect angle cue, until the target approaches a few pixels in size, at which point there is not enough information remaining to make an accurate aspect determination (stim level ~4.0).

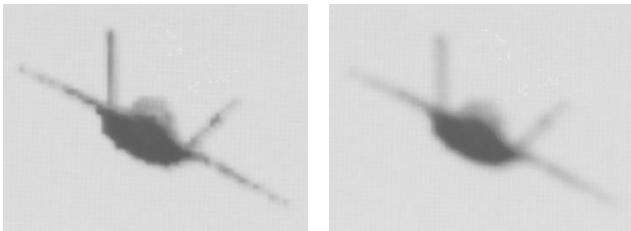


Figure 5: Illustration of the effect of bilinear warp (left) and bicubic triangular (right). Both luminance images were recorded in the center position at a stim level of 3.0 (distance of 1 km).

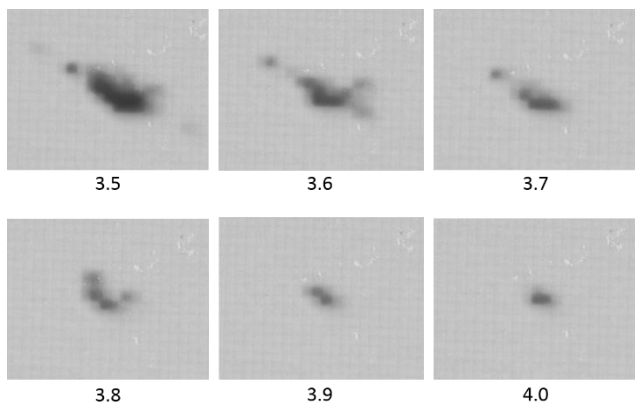


Figure 6: Model luminance at various stim levels (labeled) for bilinear warp in the central region of the projected image.

Both the body and wing contrast was measured for each target position using both bilinear warp (Figures 6 and 7) and bicubic triangular warp (Figure 8). It can be seen that the body contrast remains well above 10% for all stim levels, which means there is sufficient contrast for the subject to see the target at any size such that the number of rendered pixels becomes the limiting factor. However, it

should be noted that the contrast (at all stim levels) degrades as the target moves away from the central region and approaches the blend zone. The wings, however, experience significant, non-uniform aliasing within different regions of the screen. For bilinear warp in the central region, the wings disappear entirely beyond 6 km (i.e., stim level 3.8), while in the blend zone the wing contrast never reaches zero, but rapidly decays to about 15% beyond 4 km, which is a very difficult contrast level for human observers to see.

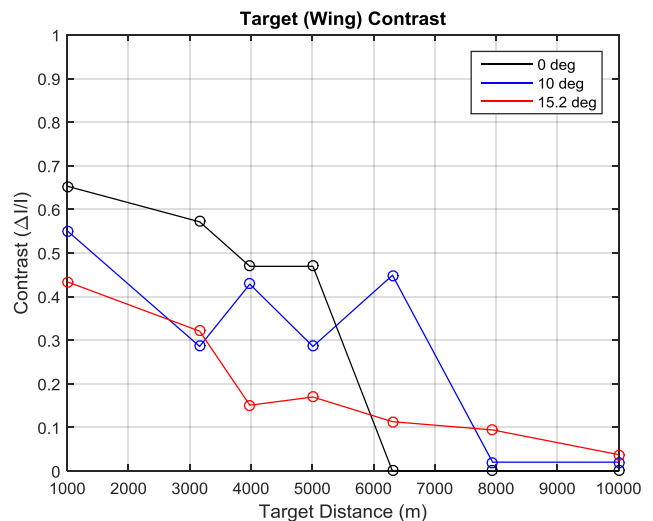
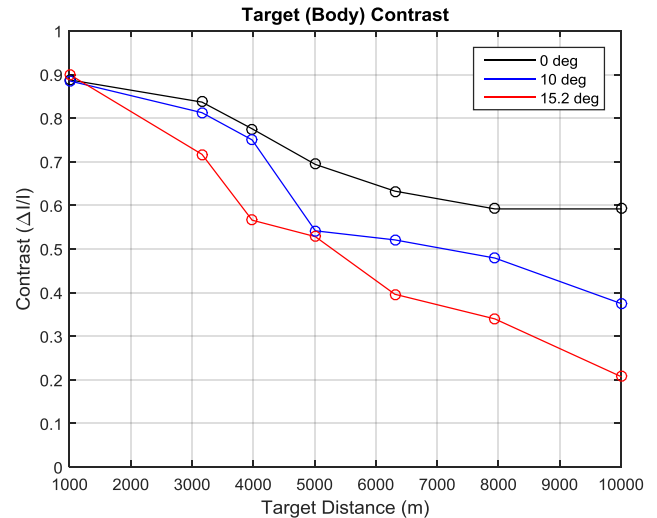


Figure 7: Measured contrast between the background and aircraft fuselage (top) and aircraft wings (bottom) for bilinear image warp.

Comparison of Figures 7 and 8 reveals that bicubic triangular resampling typically results in lower overall contrast for most target conditions. Notably, the wings are lost due to aliasing with bicubic resampling at a closer

distance than bilinear resampling methods due to the fact that bicubic interpolation provides a smoother blending (i.e., softer edges) compared to the sharper edges maintained by bilinear resampling [6].

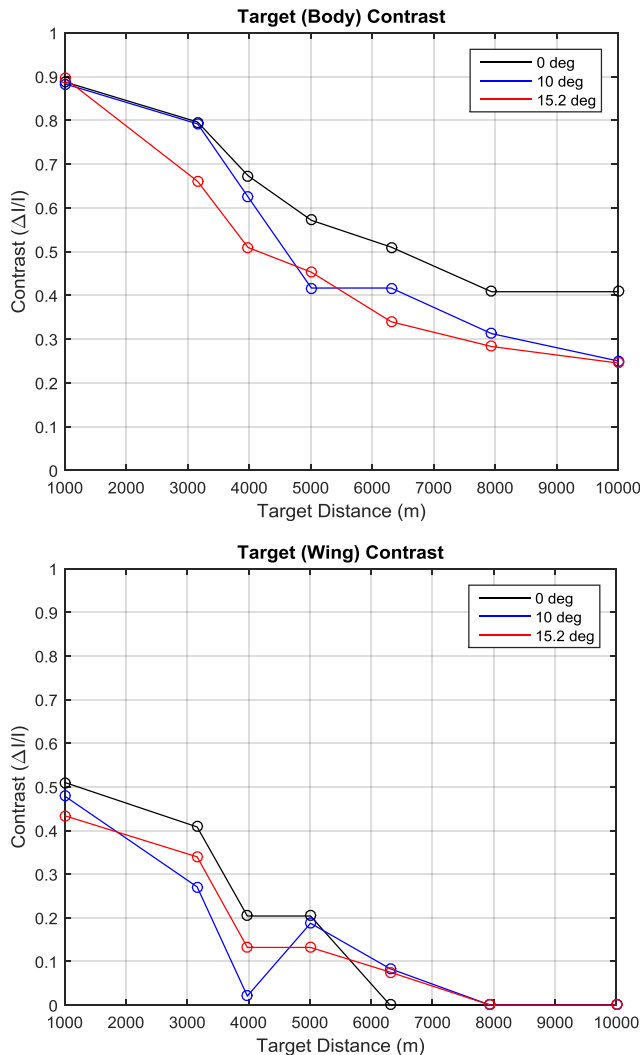


Figure 8: Measured contrast between the background and aircraft fuselage (top) and aircraft wings (bottom) for bicubic triangular image warp. The legend identifies the position of the measurement relative to the projector center.

Psychometric Measurements

The psychometric assessment consisted of 4 well-practiced observers corrected to 20/20 or better visual acuity. The display regions in which each assessment was performed, as well as the warp methods used, were presented to the subjects in a balanced order such that no two subjects proceeded through the conditions in identical order. Figure 9 illustrates the mean observation threshold versus target

location/region across the display, while Figure 10 illustrates the mean observation threshold versus warp condition. Visual inspection of Figures 9 and 10 immediately reveals that neither warp method, nor target position, yield any significant impact upon visual identification of aspect angle. All warp conditions and target positions exhibit relatively large, overlapping standard deviations, with mean fuselage identification thresholds between 1.5 to 2.0 arcminutes. This corresponds to a target width of approximately 3 to 4 pixels, or stim level of ~3.9, for all conditions.

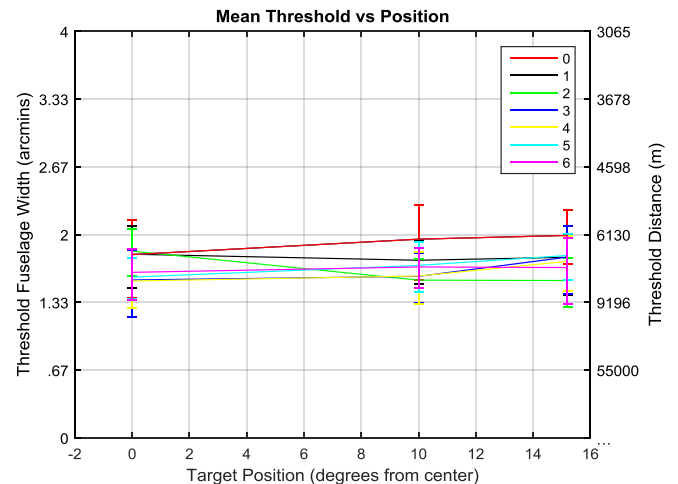


Figure 9: Mean threshold, in arcminutes, versus target position for each of seven warp conditions. Threshold stim levels range from 3.8 to 3.9 (6.3km to 8km)

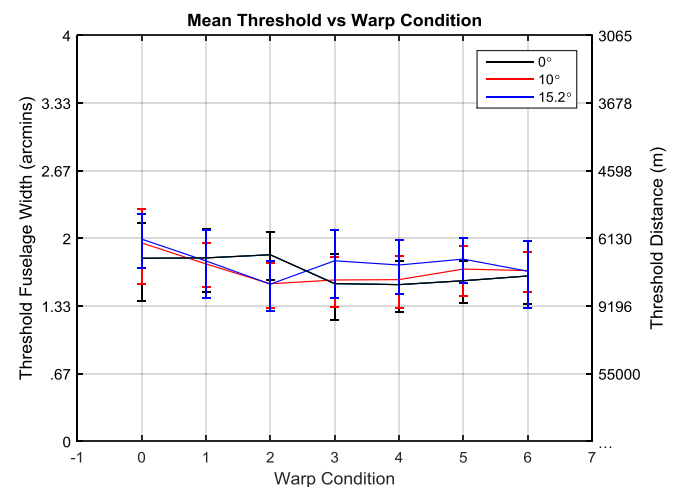


Figure 10: Mean threshold versus warp condition for each of three target positions. Threshold stim levels range from 3.8 to 3.9 (6.3km to 8km)

DISCUSSION

There are several possible explanations that may account for the lack of effect of warp method or target position on visual performance.

First, it is likely that the eye-limiting resolution of this visual system is sufficiently acute, and the aliasing errors due to geometry correction sufficiently small, that there is no performance impact for very small stimuli, as presented here. The mean target size threshold of each observer is on the order of 3 to 4 pixels and may be localized to a degree that geometry-based warp errors are not present, or have no significant effect. Thus, the nature of this stimulus may not be appropriate for measuring the impact of geometric warp on human perception

Second, it may be that the greatest degree of aliasing is introduced by the image generator used to render these targets, and not by the geometric warp. Thus, corrections to the geometric warp will not add any information that was previously removed by aliasing within the image generator. The image generator used for this experiment does not implement any form of antialiasing due to performance degradation. In previous work with much lower resolution displays, the effect of antialiasing had a small effect on the performance of a similar task, but was not consistent [7].

Third, it is possible that the geometric warp methods under investigation only impart an aesthetic improvement on the image, rather than a visual performance advantage. Inspection of Figure 5 shows that the different geometric warp methods do indeed remove much of the granularity/pixelation of small targets. Although the subjective appearance is improved, there was no corresponding performance improvement.

It is also noteworthy that visual performance did not decrease in the blend zone, which demonstrates the high fidelity of the geometry correction in the blend zone.

CONCLUSION

In this work each of the six new warp methods available in the NVIDIA Warp API was evaluated and compared to the legacy bilinear warp method. While subjective image improvements were observed, none of the new warp methods imparted any improvement in visual performance when measured by an aircraft aspect identification task. Furthermore, the placement of the target in different regions of the display, including a blended region between two displays, had no significant effect on visual performance.

ACKNOWLEDGMENTS

The use of NVIDIA or Scalable Display Systems products to conduct this work is not an endorsement of these products by the U.S. Government. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the Air Force, the Department of Defense, or the U.S. Government.

REFERENCES

- [1] J. Archdeacon, J. P. Gaska, and S. Timoner, "An operationally based vision assessment simulator for domes," in *Proceedings of the IMAGE 2012 Conference*, 2012.
- [2] J. P. Gaska, M. Winterbottom, W. Sweet, and J. Rader, "Pixel size requirements for eye-limited flight simulation," in *Proceedings of the IMAGE 2010 Conference*, 2010.
- [3] C. M. Covas, J. M. Lindholm, and C. A. Eidman, "P-28: effects of image generation variables on aircraft-roll detection range," *Dig. Tech. Pap.*, vol. 37, no. 1, pp. 299–302, 2006.
- [4] Z. Lu and B. Doshier, *Visual Psychophysics: From Laboratory to Theory*. Cambridge, MA: The MIT Press, 2013.
- [5] F. A. Kingdom and N. Prins, *Psychophysics: A Practical Introduction*. London: Academic Press, 2010.
- [6] J. A. Parker, R. V. Kenyon, and D. E. Troxel, "Comparison of interpolating methods for image resampling," *IEEE Trans. Med. Imaging*, vol. 2, no. 1, pp. 31–39, 1983.
- [7] Geri, G. and Winterbottom, M. "Effect of display resolution and antialiasing on the discrimination of simulated-aircraft orientation," *Displays*, 26, pp. 159 – 169, 2005.